

# A FAST ALGORITHM FOR THE CONVOLUTION OF FUNCTIONS WITH COMPACT SUPPORT USING FOURIER EXTENSIONS

KUAN XU\*, ANTHONY P. AUSTIN†, AND KE WEI‡

**Abstract.** We present a new algorithm for computing the convolution of two compactly supported functions. The algorithm approximates the functions to be convolved using Fourier extensions and then uses the fast Fourier transform to efficiently compute Fourier extension approximations to the pieces of the result. The complexity of the algorithm is  $\mathcal{O}(N(\log N)^2)$ , where  $N$  is the number of degrees of freedom used in each of the Fourier extensions.

**Key words.** convolution, fast Fourier transform, Fourier extension, Fredholm convolution integral, Toeplitz matrix, Volterra convolution integral

**AMS subject classifications.** 44A35, 65R10, 65T40

**1. Introduction.** Let  $f : [a, b] \rightarrow \mathbb{C}$  and  $g : [c, d] \rightarrow \mathbb{C}$  be two compactly supported functions on the real line. We consider the problem of approximating the convolution  $h = f * g$ , defined via the integral

$$(1.1) \quad h(x) = \int_{-\infty}^{\infty} f(t)g(x-t) dt = \int_{\max(a, x-d)}^{\min(b, x-c)} f(t)g(x-t) dt, \quad x \in [a+c, b+d].$$

Since convolution is a commutative operation, we lose no generality in assuming that  $g$  has the larger domain:  $d - c \geq b - a$ . In this case, the limits of integration in (1.1) can be visualized for each value of  $x$  by the diagram in Figure 1.1. Splitting  $h$  into the three pieces suggested by the diagram, we have

$$(1.2a) \quad \begin{cases} h^L(x) = \int_a^{x-c} f(t)g(x-t)dt & x \in [a+c, b+c], \\ h^M(x) = \int_a^b f(t)g(x-t)dt & x \in [b+c, a+d], \\ h^R(x) = \int_{x-d}^b f(t)g(x-t)dt & x \in [a+d, b+d]. \end{cases}$$

We refer to  $h^L$ ,  $h^M$ , and  $h^R$  as the left, middle, and right pieces of the convolution, respectively. The integrals for the left and right pieces are known as *Volterra* convo-

---

\*School of Mathematics, Statistics, and Actuarial Science, University of Kent, Canterbury, CT2 7NF, UK ([k.xu@kent.ac.uk](mailto:k.xu@kent.ac.uk)). The work of this author was supported by the Royal Society under Research Grant RG160236.

†Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Lemont, IL 60439 ([austina@anl.gov](mailto:austina@anl.gov)). The work of this author was supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.

‡Department of Mathematics, University of California, Davis, One Shields Ave., Davis, CA 95616 ([kewei@math.ucdavis.edu](mailto:kewei@math.ucdavis.edu)). The work of this author was supported by the National Science Foundation under grant number DTRA-DMS 1322393.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. <http://energy.gov/downloads/doe-public-access-plan>.




Fig. 1.1: Schematic illustrating how the limits of integration in (1.1) vary with  $x$ , assuming that  $d - c \geq b - a$ . Given a value of  $x$  on the horizontal axis, the limits range over the values of  $t$  indicated by the corresponding “slice” of the parallelogram.

lution integrals because each has a limit that varies with  $x$ , while the integral for the middle piece, which has constant limits, is a *Fredholm* convolution integral.

In [8], Hale and Townsend propose an algorithm for approximating  $h$  that first approximates  $f$  and  $g$  by polynomials represented in the Legendre basis and then convolves the approximations using a version of the convolution theorem for Legendre polynomials. The result is a piecewise polynomial representation that can be evaluated at any  $x$  in the domain of  $h$  to yield an approximation to  $h(x)$ . If the polynomials used to approximate  $f$  and  $g$  have degree at most  $M$ , their algorithm produces an approximation to  $h$  in  $\mathcal{O}(M^2)$  operations.

In this article, we present a new algorithm for approximating  $h$  using Fourier extensions instead of polynomials. The algorithm first constructs Fourier extension approximations to  $f$  and  $g$  from samples at equispaced points and then uses these to compute Fourier extension approximations to  $h^L$ ,  $h^M$ , and  $h^R$ . The advantage to using Fourier representations is that the latter computations can be carried out in just  $\mathcal{O}(N \log N)$  operations using the fast Fourier transform (FFT), where  $N$  is the number of degrees of freedom in each of the Fourier extensions for  $f$  and  $g$ . Since the extensions can be computed in just  $\mathcal{O}(N(\log N)^2)$  operations by using recently developed techniques [11], the overall complexity of our algorithm is  $\mathcal{O}(N(\log N)^2)$ .

Because Fourier extensions and polynomials approximate functions at different rates, the values of  $M$  and  $N$  defined in the preceding paragraphs needed to attain the same level of accuracy in the corresponding approximations to  $f$  and  $g$  may differ. It is therefore not immediately clear from the preceding description how our method compares with the one from [8]. We explore these issues in some of our numerical experiments below.

Our discussion is organized as follows. In Section 2, we briefly review some concepts related to Fourier extensions. In Section 3, we present the details of our algorithm. In Section 4, we test our algorithm on a variety of numerical examples, and in Section 5, we conclude the paper with some additional remarks about future work.

**2. Fourier extensions.** Fourier extensions are a technique for approximating functions on compact subintervals of the real line by using Fourier series, even if they are not periodic. Briefly, the idea is to smoothly extend the function to a periodic function on a larger domain and then approximate the extension using Fourier series.

In practice, we do not compute such an extension explicitly but, instead, directly approximate the original function using the Fourier basis from the larger domain.

In symbols, suppose that  $F$  is a complex-valued function on  $[-1, 1]$ . For fixed  $T > 1$ , we seek an approximation to  $F$  of the form

$$(2.1) \quad F(x) \approx G_n(x), \quad G_n(x) = \sum_{k=-n}^n c_k e^{ik\pi x/T}, \quad x \in [-1, 1].$$

The function  $G_n$  is a trigonometric polynomial of degree  $n$  that is periodic on the interval  $[-T, T]$ . The value  $T$  is known as the *extension parameter* and is commonly taken to be 2, although our algorithm will involve Fourier extension approximations of lengths other than this. Note that the approximation is not expected to be good for  $x \notin [-1, 1]$ , and in practice it rarely is.<sup>1</sup>

The approximation in (2.1) is usually made with respect to a least-squares criterion. One can choose  $G_n$  to minimize  $\|F - G_n\|_{L^2([-1, 1])}$ , where  $\|\cdot\|_{L^2([-1, 1])}$  is the  $L^2$  norm on  $[-1, 1]$ ; this is called the *continuous Fourier extension* problem. Alternatively, one can solve a *discrete Fourier extension problem* obtained by selecting  $M = 2m + 1 > N = 2n + 1$  points  $x_{-m}, \dots, x_m \in [-1, 1]$  and demanding that  $G_n$  minimize  $\sum_{j=-m}^m |F(x_j) - G_n(x_j)|^2$ . The number of points  $M$  is typically chosen as some fixed multiple of  $N$ , the number of degrees of freedom available in the extension. That is, we have  $M = \gamma N$  for some  $\gamma \geq 1$ , known as the *oversampling factor*. Empirical studies find that taking  $\gamma = 2$  yields adequate results in practice [2]. The most commonly used points are the equispaced points  $x_j = j/m$ , since these enable one to employ numerical methods based on the FFT. Since the FFT is integral to the success of our algorithm, we work exclusively with discrete Fourier extensions based on equispaced points in this article.

It turns out that both the continuous and discrete extension problems are poorly conditioned in that the coefficients  $c_k$  in (2.1) are highly sensitive to perturbations in  $F$  (measured in the appropriate 2-norm). Nevertheless, when solving these problems numerically, while the computed  $c_k$  may differ greatly from those of the exact solution, the function  $G_n$  that they determine can still be an excellent approximation to  $F$  provided that  $n$  is large enough that  $F$  can be resolved. These phenomena are studied carefully in [1, 2, 9], to which we refer the reader for further details. We will observe how these effects impact the stability of our proposed algorithm later in our numerical experiments.

Several methods are available for solving the discrete Fourier extension problem in equispaced points. The one we employ is the recently developed algorithm due to Matthysen and Huybrechs [11], which exploits the rapid decay of the singular values of the matrix associated with the least-squares problem using randomized techniques to compute a solution in just  $\mathcal{O}(N(\log N)^2)$  operations.<sup>2,3</sup> At the time of this writing, an implementation of their algorithm in the Julia language is publicly available on GitHub.<sup>4</sup> We have used this implementation to carry out all of the computations that we present here.

<sup>1</sup>Indeed,  $F$  may not even be defined for  $x \notin [-1, 1]$ . Moreover,  $G_n$  typically does not converge to anything outside  $[-1, 1]$  as  $n$  increases [9].

<sup>2</sup>An earlier algorithm of a similar flavor due to Lyon [10] attains the same complexity; however, it can be used to compute extensions only with  $T = 2$ .

<sup>3</sup>The implied constant in the big- $\mathcal{O}$  symbol will depend on  $T$  and  $\gamma$ . We will not address this issue here and will generally stick to the standard choices  $T = 2$  and  $\gamma = 2$ . See [11] for further information and experiments.

<sup>4</sup><https://github.com/daanhb/FrameFun.jl>

For further general information about Fourier extensions, we direct the reader to [2, 9, 11] and the references therein.

**3. Fast convolution algorithm.** The idea behind our proposed algorithm is simple. To approximate a piece (left, middle, or right) of  $h = f * g$ , we first approximate  $f$  and  $g$  using Fourier extensions. Then, we convolve these approximations to obtain a Fourier extension approximation to the result. Since the approximations are expressed in the Fourier basis, their convolutions can be computed efficiently by using the FFT. If we use extensions with  $N$  degrees of freedom to approximate  $f$  and  $g$ , then the first phase of the algorithm can be performed in  $\mathcal{O}(N(\log N)^2)$  operations by using the techniques mentioned at the end of the preceding section. The second phase takes  $\mathcal{O}(N \log N)$  operations, so the overall complexity is  $\mathcal{O}(N(\log N)^2)$ .

We now present the details of our algorithm in full.

**3.1. Changing variables.** Although Fourier extensions can be computed for functions defined on arbitrary intervals, it is convenient to describe our algorithm in terms of functions defined on basic intervals centered at the origin. To this end, we define the following mapped versions of  $f$  and  $g$ :

$$(3.1) \quad \begin{aligned} \tilde{f}(y) &= f\left(\frac{b-a}{2}y + \frac{b+a}{2}\right), & y \in [-1, 1] \\ \tilde{g}(y) &= g\left(\frac{b-a}{2}y + \frac{d+c}{2}\right), & y \in [-\rho, \rho], \end{aligned}$$

where  $\rho = (d-c)/(b-a) \geq 1$ . (Recall our assumption that  $d-c \geq b-a$ .)

One can easily relate the convolution of  $f$  and  $g$  to that of  $\tilde{f}$  and  $\tilde{g}$ . Define

$$(3.2) \quad \begin{aligned} \tilde{h}^L(y) &= \int_{-1}^y \tilde{f}(t) \tilde{g}(y - \rho - t) dt \\ \tilde{h}^R(y) &= \int_y^1 \tilde{f}(t) \tilde{g}(y + \rho - t) dt \end{aligned} \quad y \in [-1, 1].$$

Then,  $\tilde{h}^L$  and  $\tilde{h}^R$  are the left and right pieces of  $\tilde{f} * \tilde{g}$ , respectively, translated so that their domains are both  $[-1, 1]$ , and

$$(3.3) \quad \begin{aligned} h^L(x) &= \frac{b-a}{2} \tilde{h}^L\left(2\frac{x-c-b}{b-a} + 1\right), & x \in [a+c, b+c], \\ h^R(x) &= \frac{b-a}{2} \tilde{h}^R\left(2\frac{x-d-a}{b-a} - 1\right), & x \in [a+d, b+d], \end{aligned}$$

as can be seen via straightforward changes of the integration variables. For the middle piece, we set

$$(3.4) \quad \tilde{h}^M(y) = \int_{-1}^1 \tilde{f}(t) \tilde{g}(y - t) dt, \quad y \in [-(\rho-1), \rho-1],$$

and we have

$$(3.5) \quad h^M(x) = \frac{b-a}{2} \tilde{h}^M\left(\frac{2x - (a+b+c+d)}{b-a}\right), \quad x \in [b+c, a+d].$$

Note that the change-of-variables maps used to define  $\tilde{f}$  and  $\tilde{g}$  in (3.1) have identical derivatives. This feature is important because it ensures that the independent

variables of  $\tilde{f}$  and  $\tilde{g}$  are scaled relative to one another in the same way as those of the original functions  $f$  and  $g$ . Without this, the mismatch in scaling would mean that no simple relationship akin to (3.3) and (3.5) would exist between  $f * g$  and  $\tilde{f} * \tilde{g}$ .

**3.2. Approximation using Fourier extensions.** The first step in our algorithm is to approximate  $\tilde{f}$  and  $\tilde{g}$  on  $[-1, 1]$  and  $[-\rho, \rho]$ , respectively, using Fourier extensions. We must select a parameter  $T$  to use for the extensions;  $T = 2$  is the standard choice. An extension for  $\tilde{f}$  with parameter  $T$  will use the Fourier basis on  $[-T, T]$ , while one for  $\tilde{g}$  will use the Fourier basis on  $[-\rho T, \rho T]$ .

For reasons that will become clear, however, we need the period on which the extension for  $\tilde{g}$  is based to be an integer multiple of that of the extension for  $\tilde{f}$ . Therefore, instead of approximating  $\tilde{g}$  with an extension on  $[-\rho T, \rho T]$ , we use an extension on  $[-\kappa T, \kappa T]$ , where  $\kappa$  is the smallest integer such that  $\kappa \geq \rho$ . This choice ensures that the extension parameter for the approximation to  $\tilde{g}$  is as close to  $T$  as possible while still allowing the period of the extension for  $\tilde{g}$  to satisfy the required condition.

Thus, if we use Fourier extensions of lengths  $M = 2m + 1$  and  $N = 2n + 1$  to approximate  $\tilde{f}$  and  $\tilde{g}$ , respectively, our approximations take the form

$$(3.6) \quad \tilde{f}(y) \approx \sum_{j=-m}^m a_j e^{ij\pi y/T}, \quad \tilde{g}(y) \approx \sum_{k=-n}^n b_k e^{ik\pi y/(\kappa T)}.$$

**3.3. Computing the left and right pieces.** We now describe how to compute  $\tilde{h}^L$  and  $\tilde{h}^R$ , from which  $h^L$  and  $h^R$  can be recovered using (3.3). We will focus on the computation of  $\tilde{h}^L$ ; the procedure for computing  $\tilde{h}^R$  is similar.<sup>5</sup> Substituting the approximations (3.6) into (3.2) and using the fact that

$$\int_{-1}^y e^{i(\kappa j - k)\pi t/(\kappa T)} dt = \begin{cases} \frac{\kappa T}{i\pi} \left( \frac{e^{i\pi(\kappa j - k)y/(\kappa T)}}{\kappa j - k} + \frac{e^{i\pi(k - \kappa j)/(\kappa T)}}{k - \kappa j} \right) & k \neq \kappa j \\ y + 1 & k = \kappa j, \end{cases}$$

we obtain

$$(3.7) \quad \begin{aligned} \tilde{h}^L(y) &\approx \sum_{j=-m}^m a_j \sum_{k=-n}^n b_k e^{ik\pi(y-\rho)/(\kappa T)} \int_{-1}^y e^{i(\kappa j - k)\pi t/(\kappa T)} dt \\ &= \sum_{j=-m}^m \tilde{a}_j e^{ij\pi y/T} + \sum_{k=-n}^n \tilde{b}_k e^{ik\pi y/(\kappa T)} + (y+1) \sum_{\substack{j=-m \\ -n \leq \kappa j \leq n}} c_j e^{ij\pi y/(\kappa T)}, \end{aligned}$$

where

$$(3.8) \quad \tilde{a}_j = \frac{\kappa T}{i\pi} a_j p_j, \quad p_j = \sum_{\substack{k=-n \\ k \neq \kappa j}}^n b_k e^{-ik\pi\rho/(\kappa T)} \frac{1}{\kappa j - k}, \quad -m \leq j \leq m,$$

and

$$(3.9) \quad \tilde{b}_k = \frac{\kappa T}{i\pi} b_k q_k e^{ik\pi(1-\rho)/(\kappa T)}, \quad q_k = \sum_{\substack{j=-m \\ \kappa j \neq k}}^m a_j e^{-ij\pi/T} \frac{1}{k - \kappa j}, \quad -n \leq k \leq n,$$

<sup>5</sup>In more detail,  $\tilde{h}^R(y)$ ,  $y \in [-1, 1]$ , is just the evaluation at  $-y$  of the left piece of  $\tilde{f}^* * \tilde{g}^*$ , where  $\tilde{f}^*(y) = \tilde{f}(-y)$  and  $\tilde{g}^*(y) = \tilde{g}(-y)$ .

and

$$(3.10) \quad c_j = a_j b_{\kappa j} e^{-ij\pi\rho/T}, \quad -m \leq j \leq m, -n \leq \kappa j \leq n.$$

The key observation to make is that the sums defining  $p_j$  and  $q_k$  in (3.8) and (3.9) can be interpreted as matrix-vector products in which the matrices involved have Toeplitz structure. Let  $b = (b_{-n}e^{-i(-n)\pi\rho/(\kappa T)}, \dots, b_n e^{-in\pi\rho/(\kappa T)})$ , a (column) vector of length  $N$ , and let  $P$  be the  $(2\kappa m + 1) \times N$  matrix whose  $(\nu, k)$  entry is

$$P_{\nu,k} = \begin{cases} \frac{1}{\nu-k} & \nu \neq k \\ 0 & \nu = k \end{cases}$$

for  $-\kappa m \leq \nu \leq \kappa m$  and  $-n \leq k \leq n$ . Let  $\hat{p} = Pb$ . If  $\hat{p}_\nu$  denotes the  $\nu$ th component of  $\hat{p}$ ,  $-\kappa m \leq \nu \leq \kappa m$ , then from (3.8) we have  $p_j = \hat{p}_{\kappa j}$ . Thus, we can compute the  $p_j$  by computing  $\hat{p}$  and keeping every  $\kappa$ th component. Similarly, if  $\hat{a}$  is the vector of length  $2\kappa m + 1$  defined by

$$\hat{a}_\nu = \begin{cases} a_{\nu/\kappa} e^{-i(\nu/\kappa)\pi/T} & \nu \equiv 0 \pmod{\kappa} \\ 0 & \text{otherwise} \end{cases}$$

for  $-\kappa m \leq \nu \leq \kappa m$  and  $q = (q_{-n}, \dots, q_n)$ , then by (3.9) we have  $q = Q\hat{a}$ , where  $Q$  is the  $N \times (2\kappa m + 1)$  matrix whose  $(k, \nu)$  entry is

$$Q_{k,\nu} = \begin{cases} \frac{1}{k-\nu} & k \neq \nu \\ 0 & k = \nu \end{cases}$$

for  $-n \leq k \leq n$  and  $-\kappa m \leq \nu \leq \kappa m$ .

Since  $P_{\nu,k}$  and  $Q_{k,\nu}$  depend only on the differences  $\nu - k$  and  $k - \nu$ , respectively, and not on  $\nu$  and  $k$  individually,  $P$  and  $Q$  are Toeplitz. Therefore, the multiplications needed to compute  $\hat{p}$  and  $q$  (and, hence, *all* the  $p_j$  and  $q_k$ ) can be carried out in  $\mathcal{O}((2\kappa m + N) \log(2\kappa m + N))$  operations by the standard technique of embedding them in circulant matrices and using the FFT [15, §4.8.4]. Since these multiplications constitute the dominant cost in computing the coefficients  $\tilde{a}_j$ ,  $\tilde{b}_k$ , and  $c_j$ , the approximation (3.7) to  $\tilde{h}^L$  can be formed in  $\mathcal{O}((2\kappa m + N) \log(2\kappa m + N))$  operations.

Noting that the first sum may be rewritten to use complex exponentials of period  $\kappa T$  by inserting factors of  $\kappa$  into the numerator and denominator of the exponent, we see that the expansion (3.7) is nearly a Fourier extension approximation to  $\tilde{h}^L$  with extension parameter  $\kappa T$ . The only problem is the appearance of the factor of  $y + 1$  multiplying the third sum. We can eliminate this factor by approximating it on  $[-1, 1]$  with a Fourier extension

$$y + 1 \approx \sum_{r=-d}^d \gamma_r e^{ir\pi y/T},$$

where the degree  $d$  is chosen large enough to attain a desired level of accuracy. Expressing this sum using complex exponentials of period  $\kappa T$  and then multiplying the result with the third sum in (3.7), we obtain an approximation

$$(3.11) \quad (y + 1) \sum_{\substack{j=-m \\ -n \leq \kappa j \leq n}}^m c_j e^{ij\pi y/(\kappa T)} \approx \sum_{\ell=-d'}^{d'} \tilde{c}_\ell e^{i\ell\pi y/(\kappa T)},$$

where  $d' = \kappa d + \min(m, \lceil n/\kappa \rceil)$  and  $\lceil x \rceil$  denotes the least integer greater than or equal to  $x$ . The coefficients  $\tilde{c}_\ell$  can be computed efficiently in  $O(d' \log(d'))$  operations by using the standard FFT-based algorithm for multiplying trigonometric polynomials.

For the usual choice of  $T = 2$ , we find experimentally that an extension with  $d = 24$  (49 basis functions) is sufficient to approximate  $y + 1$  to machine precision uniformly on  $[-1, 1]$  in double-precision arithmetic. Note that this approximation can be computed once, offline, and then reused as many times as necessary.

**3.4. Computing the middle piece.** The middle piece is handled similarly to the left and right pieces. In fact, it is slightly simpler because the limits of integration in (3.4) are constant, whereas those in (3.2) vary with  $y$ . Inserting the approximations (3.6) into (3.4) and noting that

$$\int_{-1}^1 e^{i(\kappa j - k)\pi t / (\kappa T)} dt = \begin{cases} \frac{2\kappa T}{\pi} \frac{\sin(\frac{\pi}{\kappa T}(\kappa j - k))}{\kappa j - k} & k \neq \kappa j \\ 2 & k = \kappa j, \end{cases}$$

we find

$$\begin{aligned} \tilde{h}^M(y) &\approx \sum_{j=-m}^m a_j \sum_{k=-n}^n b_k e^{ik\pi y / (\kappa T)} \int_{-1}^1 e^{i(\kappa j - k)\pi t / (\kappa T)} dt \\ (3.12) \quad &= \sum_{k=-n}^n \tilde{b}_k e^{ik\pi y / (\kappa T)} + \sum_{\substack{j=-m \\ -n \leq \kappa j \leq n}}^m \tilde{c}_j e^{i\kappa j \pi y / (\kappa T)}, \end{aligned}$$

where

$$(3.13) \quad \tilde{b}_k = \frac{2\kappa T}{\pi} b_k r_k, \quad r_k = \sum_{\substack{j=-m \\ \kappa j \neq k}}^m a_j \frac{\sin(\frac{\pi}{\kappa T}(\kappa j - k))}{\kappa j - k}, \quad -n \leq k \leq n,$$

and

$$(3.14) \quad \tilde{c}_j = 2a_j b_{\kappa j}, \quad -m \leq j \leq m, -n \leq \kappa j \leq n.$$

The expansion (3.12) is a Fourier extension approximation to  $\tilde{h}^M$  with period  $\kappa T$ . The dominant cost in computing the coefficients  $\tilde{b}_k$  and  $\tilde{c}_j$  in this expansion comes from computing the values  $r_k$  defined in (3.13); however, as was the case for  $p_j$  and  $q_k$  in (3.8) and (3.9), we can accelerate the computation of the  $r_k$  by taking advantage of the structure of the sum. Letting  $r = (r_{-n}, \dots, r_n)$ , a vector of length  $N = 2n + 1$ , we have  $r = R\hat{a}$ , where  $\hat{a} = (\hat{a}_{-\kappa m}, \dots, \hat{a}_{\kappa m})$  is the vector of length  $2\kappa m + 1$  defined by

$$\hat{a}_\nu = \begin{cases} a_{\nu/\kappa} & \nu \equiv 0 \pmod{\kappa} \\ 0 & \text{otherwise} \end{cases}$$

for  $-\kappa m \leq \nu \leq \kappa m$ , and  $R$  is the  $N \times (2\kappa m + 1)$  Toeplitz matrix with entries

$$R_{k,\nu} = \begin{cases} \frac{\sin(\frac{\pi}{\kappa T}(\nu - k))}{\nu - k} & k \neq \nu \\ 0 & k = \nu \end{cases}$$

for  $-n \leq k \leq n$  and  $-\kappa m \leq \nu \leq \kappa m$ . As before, this matrix can be applied to  $\hat{a}$  in  $\mathcal{O}((2\kappa m + N) \log(2\kappa m + N))$  operations by using the FFT.

**3.5. Summary of the main algorithm.** We summarize our algorithm for computing a Fourier extension approximation to (1.1) as follows:

1. Choose an extension parameter  $T$  on which to base the approximations; taking  $T = 2$  is standard and should suffice in general.
2. Compute Fourier extension approximations to  $\tilde{f}$  and  $\tilde{g}$  as prescribed by (3.6) using one of the recently developed fast algorithms mentioned in Section 2.
3. Compute Fourier extension approximations to  $\tilde{h}^L$ ,  $\tilde{h}^M$ , and  $\tilde{h}^R$  by appropriately convolving the Fourier extensions computed in Step 2. This can be done efficiently by using the FFT to compute the coefficients of the approximations as governed by (3.7) through (3.11) for the left piece (and similarly for the right piece) and (3.12) through (3.14) for the middle piece.
4. Undo the changes of variables according to (3.3) and (3.5) to obtain Fourier extension approximations to  $h^L$ ,  $h^M$ , and  $h^R$ .

**3.6. An alternative approach to the left and right pieces.** In some applications, only the left and/or right piece of the convolution integral is of interest. When this is the case, one may be able to save some operations by approximating only the portion of  $g$  that enters into the convolution instead of all of  $g$ . We now describe how to do this. As in Section 3.3, we focus on the computation of the left piece; the right piece can be handled analogously.

Let

$$\tilde{g}^L(y) = g\left(\frac{b-a}{2}y + \frac{b-a}{2} + c\right), \quad y \in [-1, 1].$$

The function  $\tilde{g}^L$  maps the portion of  $g$  defined on  $[c, c + (b-a)]$  to  $[-1, 1]$ ; this is precisely the portion of  $g$  that is involved in the left piece of  $f * g$ . One can easily see that

$$(3.15) \quad \tilde{h}^L(y) = \int_{-1}^y \tilde{f}(t) \tilde{g}^L(y-1-t) dt, \quad y \in [-1, 1],$$

where  $\tilde{h}^L$  is as defined in (3.2). To approximate  $\tilde{h}^L$  using (3.15), we approximate  $\tilde{f}$  and  $\tilde{g}^L$  using Fourier extensions:

$$(3.16) \quad \tilde{f}(y) \approx \sum_{j=-m}^m a_j e^{ij\pi y/T}, \quad \tilde{g}^L(y) \approx \sum_{k=-n}^n b_k e^{ik\pi y/T}.$$

The difference between (3.6) and (3.16) is that since the domains of  $\tilde{f}$  and  $\tilde{g}^L$  are the same, they can be approximated by extensions with the same parameter using complex exponentials with the same period as if  $\kappa$  were equal to 1. Doing this will, in general, reduce the length of the FFTs required to compute the coefficients of the Fourier extension approximation to the convolution. Moreover, since  $\tilde{g}^L$  represents only a part of  $g$ , one probably will be able to get away with using a shorter series (smaller value of  $n$ ) to obtain an approximation of a given accuracy than one would need in order to approximate all of  $g$  via  $\tilde{g}$ .

From here, we proceed exactly as in Section 3.3, replacing  $\tilde{g}$  with  $\tilde{g}^L$  and taking  $\kappa = 1$  in all formulae. For completeness, we give the simplified formulae in full here.



By inserting (3.16) into (3.15), we obtain

$$\begin{aligned}
 \tilde{h}^L(y) &\approx \sum_{j=-m}^m a_j \sum_{k=-n}^n b_k e^{ik\pi(y-1)/T} \int_{-1}^y e^{i\pi(j-k)t/T} dt \\
 (3.17) \quad &= \sum_{j=-m}^m \tilde{a}_j e^{ij\pi y/T} + \sum_{k=-n}^n \tilde{b}_k e^{ik\pi y/T} + (y+1) \sum_{\ell=-\min(m,n)}^{\min(m,n)} c_\ell e^{i\ell\pi y/T},
 \end{aligned}$$

where

$$(3.18) \quad \tilde{a}_j = \frac{T}{i\pi} a_j p_j e^{-ij\pi/T}, \quad p_j = \sum_{\substack{k=-n \\ k \neq j}}^n b_k \frac{e^{i\pi(j-k)/T}}{j-k}, \quad -m \leq j \leq m,$$

and

$$(3.19) \quad \tilde{b}_k = \frac{T}{i\pi} b_k q_k e^{-ik\pi/T}, \quad q_k = \sum_{\substack{j=-m \\ j \neq k}}^m a_j \frac{e^{i\pi(k-j)/T}}{k-j}, \quad -n \leq k \leq n,$$

and

$$(3.20) \quad c_\ell = a_\ell b_\ell e^{-i\ell\pi/T}, \quad -\min(m,n) \leq \ell \leq \min(m,n).$$

From (3.18), we have  $p = Pb$ , where  $p = (p_{-m}, \dots, p_m)$  and  $b = (b_{-n}, \dots, b_n)$  are vectors of length  $M = 2m + 1$  and  $N = 2n + 1$ , respectively, and  $P$  is the  $M \times N$  Toeplitz matrix whose  $(j, k)$  entry is

$$P_{j,k} = \begin{cases} \frac{e^{i\pi(j-k)/T}}{j-k} & j \neq k \\ 0 & j = k \end{cases}$$

for  $-m \leq j \leq m$  and  $-n \leq k \leq n$ . Similarly, from (3.19), we have  $q = Qa$ , where  $q = (q_{-n}, \dots, q_n)$ ,  $a = (a_{-m}, \dots, a_m)$ , and  $Q$  is the  $N \times M$  Toeplitz matrix defined by

$$Q_{k,j} = \begin{cases} \frac{e^{i\pi(k-j)/T}}{k-j} & j \neq k \\ 0 & j = k \end{cases}$$

for  $-m \leq j \leq m$  and  $-n \leq k \leq n$ .

This takes care of the first and second sums in (3.17). The multiplication of the third sum by  $y+1$  can be handled in exactly the same manner as described in Section 3.3.

**4. Numerical examples.** We now examine the characteristics of the proposed algorithm and evaluate its performance on a few numerical examples. The Fourier extensions were computed by using the publicly available Julia implementation of the Matthysen-Huybrechts algorithm mentioned in Section 2. In all examples, we use an extension parameter of  $T = 2$  and an oversampling factor of  $\gamma = 2$  unless otherwise specified.




Fig. 4.1: Convolution of two boxes. (a) The functions  $f$  and  $g$  being convolved and the approximation to their convolution  $h$ . (b) The error in the approximation. Because of the simplicity of this example, the error comes almost entirely from the precomputed approximation to  $y + 1$  in (3.11).

**4.1. Convolution of two boxes.** As an easy first test, we set  $f$  and  $g$  to be the characteristic functions of the intervals  $[-1, 1]$  and  $[-2, 2]$ , respectively. Figure 4.1a depicts  $f$ ,  $g$ , and the computed approximation to their (full) convolution. Figure 4.1b displays the pointwise absolute error in the computation.

Since  $f$  and  $g$  can each be represented exactly by using just a single (zero-frequency) Fourier coefficient, this example is perhaps the simplest possible problem to which our algorithm may be applied. The error in the result is due almost entirely to the error in the precomputed approximation to  $y + 1$  in (3.11) needed for the left and right pieces.

**4.2. A convolution from renewal theory.** For our next example, we consider the Volterra convolution integral equation

$$(4.1) \quad f(x) = g(x) + \int_0^x f(t)g(x-t) dt, \quad x \geq 0,$$

with

$$g(x) = \frac{1}{2}x^2e^{-x}.$$

This equation arises in renewal theory [5, Example 1.4.3], [7], and one can show using Laplace transforms that the solution is

$$f(x) = \frac{1}{3} - \frac{1}{3} \left( \cos \frac{\sqrt{3}}{2}x + \sqrt{3} \sin \frac{\sqrt{3}}{2}x \right) e^{-3x/2}.$$

We verify this solution by using our algorithm to approximate the integral that appears in (4.1) and comparing the result with  $f - g$ .

We approximate  $f$  and  $g$  (which are similar but not identical) on  $[0, 1]$  using Fourier extensions of length 71. Since only the left piece of  $f * g$  is desired, we use the simplified version of our algorithm given in Section 3.6 to perform the computation. Figure 4.2a depicts  $f$ ,  $g$ , and the computed approximation to the left piece  $h^L$  of




Fig. 4.2: A convolution related to renewal theory. (a) The functions  $f$  and  $g$  being convolved and the approximation to the left piece of their convolution. (b) Absolute error in the approximations to  $f$ ,  $g$ , and the left piece of their convolution.

their convolution, which is also defined on  $[0, 1]$ . The pointwise absolute error in the approximations to  $f$ ,  $g$ , and  $h^L$  is displayed in Figure 4.2b.

The functions  $f$  and  $g$  are not complicated and do not require many coefficients to represent accurately, making this still a relatively easy example. Nevertheless, the computation illustrates an interesting fact about approximating convolutions using Fourier extensions. We observe from Figure 4.2b that while the error in the Fourier extension approximations to  $f$  and  $g$  is as high as  $10^{-13}$ , the error in the approximation to their convolution is uniform at a level just under  $10^{-16}$ . That is, the accuracy in the approximation to  $h^L$  (measured in the uniform norm) is roughly three orders of magnitude greater than the accuracy in the approximations used to compute it. This phenomenon can be partially explained by taking into account the fact that  $h^L$  is roughly 10 times smaller than  $f$  and  $g$  (i.e., we should consider measuring the error in a relative rather than an absolute sense), but this reasoning is not enough to account for everything. The other key observation is that the regions of comparatively large error in the approximations to  $f$  and  $g$  are localized to the endpoints, so they do not contribute much to the convolution integral.

One can easily make this reasoning precise. Suppose that we have approximated  $f$  and  $g$  by  $\tilde{f} = f + \delta_f$  and  $\tilde{g} = g + \delta_g$ , respectively, where  $\delta_f$  and  $\delta_g$  are functions that satisfy  $\|\delta_f\|_\infty \leq \varepsilon$  and  $\|\delta_g\|_\infty \leq \varepsilon$  for some  $\varepsilon > 0$ . Here,  $\|\cdot\|_\infty$  denotes the uniform norm. The error in the approximation to  $f * g$  is

$$(\tilde{f} * \tilde{g})(x) - (f * g)(x) = (\delta_f * g)(x) + (f * \delta_g)(x) + (\delta_f * \delta_g)(x).$$

Noting that the third term on the right-hand side is  $\mathcal{O}(\varepsilon^2)$  as  $\varepsilon \rightarrow 0$ , upon taking absolute values and applying Hölder's inequality to the integrals that define the remaining two terms, we obtain

$$(4.2) \quad \|\tilde{f} * \tilde{g} - f * g\|_\infty \leq \|g\|_\infty \|\delta_f\|_1 + \|f\|_\infty \|\delta_g\|_1 + \mathcal{O}(\varepsilon^2),$$

where  $\|\cdot\|_1$  denotes the  $L^1$ -norm over the appropriate domain.

From (4.2), we see that the uniform-norm error in the convolution can be bounded by an expression in which the significant terms depend on  $\|\delta_f\|_1$  and  $\|\delta_g\|_1$  but not




Fig. 4.3: Results for the oscillatory convolution of Section 4.3. (a) Maximum error in the convolution on a grid of 8192 equispaced points in  $[-2, 2]$  versus the number of degrees of freedom  $N$  in the Fourier extensions. (b) Time required to compute the convolution versus  $N$ . The blue and green lines show the time for approximating  $f$  and  $g$  and for the subsequent computations required to convolve the approximations, respectively. The red line shows the total time taken (the sum of the two).

$\|\delta_f\|_\infty$  and  $\|\delta_g\|_\infty$ . That is, the *integrals* of the errors in the approximations to  $f$  and  $g$  determine the error in the approximation to the result. Hence, these errors must be large over a significant portion of the domain before their impacts become important.

In Figure 4.2b, the areas of large error in  $f$  and  $g$  near the endpoints together occupy well less than  $1/10$  of the total interval, and they turn out to be relatively insignificant. This is useful to know, since error profiles similar to those in Figure 4.2b are not uncommon in practice.

**4.3. Convolution of oscillatory functions.** Next, we apply our algorithm to approximate the (full) convolution of the oscillatory functions  $f(x) = \sin(100x) + x/50$  and  $g(x) = \cos(200x)^2$ , both defined on  $[-1, 1]$ . The exact answer can be found by using a symbolic computing system or a table to evaluate the necessary integrals. We approximate  $f$  and  $g$  using Fourier extensions with  $N$  degrees of freedom for several values of  $N$  ranging from 5 to 605 and measure both the time taken to compute the convolution (averaged over 3 trials) and the maximum absolute error in the result on a grid of 8192 equispaced points in  $[-2, 2]$ .

The results are displayed in Figure 4.3. Figure 4.3a shows that the approximation to the convolution has converged to the exact solution to nearly machine precision on  $[-2, 2]$  once  $N$  exceeds approximately 525. Figure 4.3b shows that the amount of time taken does indeed appear to grow at a rate bounded by  $O(N(\log N)^2)$  as predicted. Moreover, the bulk of the time is clearly spent building the Fourier extension approximations to  $f$  and  $g$ . This is expected because the singular value decomposition computed by the Matthysen-Huybrechts algorithm to solve the extension problems is considerably more expensive than the handful of FFTs needed to convolve the extensions, even taking the randomized nature of the algorithm into account.

Figure 4.3a displays a trend that initially appears alarming. For several values of  $N$ , the error in the approximation is exceedingly high—on the order of  $10^9$ . The behavior disappears around  $N = 125$ , at which point the error settles down to a nominal level of  $10^{-3}$  before beginning its eventual decay to machine precision around

$N = 500$ . Looking more closely, we find that the instability arises as a result of the coefficients of the Fourier extension approximations to  $f$  and  $g$  being excessively large, with some as high as  $10^{11}$  for the values of  $N$  for which the instability is most severe.

This behavior is actually a well-understood property of Fourier extensions; for extensive studies with quantitative results, we direct the reader to [1, 2]. The key point is that it (provably) does not persist indefinitely as  $N$  increases; it emerges only in a regime in which the function being approximated has not been adequately resolved. Thus, while this effect means that our algorithm exhibits instability for some values of  $N$ , the instability is harmless, since those values of  $N$  all correspond to extensions with insufficiently many degrees of freedom to approximate the function accurately, meaning that a stably computed approximation would be of little utility anyway.

**4.4. Comparison with methods based on polynomials.** For our final experiment, we compare our algorithm to the one of Hale and Townsend [8], which approximates  $f$  and  $g$  by polynomials instead of Fourier extensions. Recall from the introduction that this method scales as  $\mathcal{O}(M^2)$ , where  $M$  is the degree of the polynomials employed. At first glance, it would appear that the method of this article, which scales as  $\mathcal{O}(N(\log N)^2)$  (where, as usual,  $N$  denotes the number of degrees of freedom in each of the Fourier extensions used), should be superior.

The problem with this reasoning is that it is not clear how  $M$  relates to  $N$ , so the complexities cannot be directly compared. Fourier extensions and polynomials approximate functions at different rates. If the value of  $N$  needed for Fourier extension approximations to  $f$  and  $g$  to attain a given accuracy is much greater than the value of  $M$  needed for polynomial approximations to attain same accuracy (e.g., if  $N \propto M^2$ ), then the polynomial-based method may win out in the end.

Attempting to provide a definitive answer to the question of how  $M$  and  $N$  relate in general is beyond the scope of this article. Nevertheless, we can show numerically that there are situations in which our method may be competitive. For  $x \in [-1, 1]$ , let  $f(x) = 1 + 1/(1 + \omega x^2)$ , a Runge-like function, where  $\omega \in \mathbb{R}$  is a given parameter. We consider the convolution  $h = f * f$ . As was the case with the previous example, the exact answer to this problem may be found by using a symbolic computing system or consulting a table of integrals.

For a range of values of  $\omega$ , we determine  $M$  such that  $f$  is approximated uniformly on  $[-1, 1]$  to within an accuracy of approximately  $10^{-10}$  by its degree- $M$  polynomial interpolant in a grid of (second-kind) Chebyshev points. We convolve this interpolant with itself using the algorithm from [8] and measure the time taken. We do not include the time taken to build the interpolant. We compare this with the time taken to approximate and convolve  $f$  with itself using our algorithm with a Fourier extension of length  $N$  that resolves  $f$  to approximately the same level of accuracy.<sup>6</sup>

As  $\omega$  increases, the spike in  $f$  near the origin becomes increasingly narrow, and  $M$  and  $N$  will need to increase accordingly; thus, this experiment should provide insight into how the different methods scale as the problem complexity increases. All our computations involving polynomials were carried out in MATLAB R2016b using the Chebfun software package [6], which contains an implementation of the algorithm

---

<sup>6</sup>While  $M$  can be determined by using features of Chebfun, at the time of this writing, construction of an efficient general procedure for determining  $N$  is a matter of ongoing study. We determined  $N$  using a brute-force search. The reported times do not include the time taken to compute  $M$  and  $N$ .




Fig. 4.4: Results for the experiment of Section 4.4. (a) Number of degrees of freedom  $N$  in the Fourier extension and the degree  $M$  of the polynomial used to approximate  $f$  to the desired level of accuracy for a given value of  $\omega$ . (b) Ratio of  $N$  to  $M$ . (c) Times required to compute the convolution using both methods. (d) Approximate uniform-norm error in the convolution computed using both methods.

from [8]. Accuracy of both the Fourier extension and polynomial approximations was measured by computing the maximum absolute error on a fine grid in  $[-1, 1]$ .

The results are displayed in Figure 4.4. Figures 4.4a and 4.4b show how the quantities  $M$  and  $N$  described in the previous paragraph vary with  $\omega$ . We observe from Figure 4.4b that  $N \approx 1.4M$  for this problem, which suggests that as  $\omega$  increases, the algorithm based on Fourier extensions should eventually exhibit superior performance. This is borne out by the plots of Figure 4.4c. Because the algorithms have been implemented in different languages using different software packages, making a fair direct comparison between the times is difficult. Nevertheless, the figure clearly shows that the Fourier extension method becomes faster once  $\omega$  is large enough. Figure 4.4d shows that all computations evaluate the convolutions to comparable levels of accuracy.

Actually, the performance of the Fourier extension method for this problem is even better than Figure 4.4c might suggest. The reason is that our implementation of the algorithm does not take advantage of the fact that  $f$  is being convolved with itself instead of with some other arbitrary function  $g$ . We therefore actually construct

the Fourier extension approximation to  $f$  *twice* over the course of the computation. Since constructing the Fourier extension is the most expensive part of the algorithm (see Figure 4.3b), taking advantage of this would lead to significant savings. The polynomial-based method can benefit similarly; however, we already exclude the times required to build the polynomial interpolants from our measurements because they are negligible, so they do not appear in Figure 4.4c as is.

We close this section by noting a crucial difference between our method and the polynomial-based method just employed: the latter bases its approximations on samples of  $f$  at Chebyshev points, whereas the former uses equispaced points. It is not uncommon for an application to demand samples be taken on an equispaced grid. If this is the case, polynomial interpolants cannot be used because of the Runge phenomenon [14, Ch. 13].

One still can use polynomials by performing a discrete least-squares approximation instead (as is done for Fourier extensions); however, this comes with some disadvantages. In order to ensure stable convergence, the number of samples taken must be proportional to the square of the degree of the polynomial used [3, 4, 13]. This severely limits the size of the degree that can be used in practice.<sup>7</sup> At the same time, the convergence rate degrades; if the scheme is stable, it can converge only at a rate that is subgeometric in the degree, even when the function being approximated is entire [12]. These effects severely blunt the general advantage polynomials enjoy over Fourier extensions with respect to rates of approximation, meaning that our method will be even more competitive in this context.

**5. Conclusion.** We have presented a new numerical method for convolving two functions with compact support. The method approximates the functions to be convolved using Fourier extensions and then convolves the extensions efficiently with the aid of the FFT to produce a Fourier extension approximation to the convolution. When the extensions are computed using one of the recently developed fast algorithms, the method takes  $\mathcal{O}(N(\log N)^2)$  operations, where  $N$  is the number of degrees of freedom in each of the Fourier extensions.

While we have worked exclusively with Fourier extensions in this article, we note that the techniques described in Section 3 can be used with any scheme that produces sum-of-exponentials approximations of the form (3.6). We plan to investigate the utility of our approach for computing convolutions in higher dimensions as well as the possibility of using discretizations based on our method for solving convolution integral equations.

**Acknowledgments.** We thank Daan Huybrechs for his patient and detailed answers to our many questions about Fourier extensions and for providing us assistance with using the FrameFun.jl software package to perform our numerical experiments. We also thank Alex Townsend for his comments on a draft of our manuscript, which led us to make several improvements.

## REFERENCES

- [1] B. ADCOCK AND D. HUYBRECHS, *Frames and numerical approximation*. In preparation, 2016.

<sup>7</sup>Assuming the constant of proportionality is 1, even storing the (dense)  $10^6 \times 10^3$  matrix for the least-squares problem with a degree-1000 polynomial requires 8 GB of memory in double precision. In contrast, a Fourier extension approximation with *twice* as many degrees of freedom created using the standard oversampling factor  $\gamma = 2$  requires only 64 MB.

- [2] B. ADCOCK, D. HUYBRECHS, AND J. MARTÍN-VAQUERO, *On the numerical stability of Fourier extensions*, Found. Comput. Math., 14 (2014), pp. 635–687.
- [3] B. ADCOCK, R. PLATTE, AND A. SHADRIN, *Optimal sampling rates for approximating analytic functions from pointwise samples*, arXiv, (2015). [arXiv:1610.04769v1 \[math.NA\]](#).
- [4] J. P. BOYD AND F. XU, *Divergence (Runge Phenomenon) for least-squares polynomial approximation on an equispaced grid and Mock-Chebyshev subset interpolation*, Appl. Math. Comput., 210 (2009), pp. 158–168.
- [5] H. BRUNNER AND P. J. VAN DER HOUWEN, *The Numerical Solution of Volterra Equations*, Elsevier, 1986.
- [6] T. A. DRISCOLL, N. HALE, AND L. N. TREFETHEN, eds., *Chebfun Guide*, Pafnuty Publications, Oxford, 2014.
- [7] W. FELLER, *On the integral equation of renewal theory*, Ann. Math. Statist., 12 (1941), pp. 243–267.
- [8] N. HALE AND A. TOWNSEND, *An algorithm for the convolution of Legendre series*, SIAM J. Sci. Comput., 36 (2014), pp. A1207–A1220.
- [9] D. HUYBRECHS, *On the Fourier extension of nonperiodic functions*, SIAM J. Numer. Anal., 47 (2010), pp. 4326–4355.
- [10] M. LYON, *A fast algorithm for Fourier continuation*, SIAM J. Sci. Comput., 33 (2011), pp. 3241–3260.
- [11] R. MATTHYSEN AND D. HUYBRECHS, *Fast algorithms for the computation of Fourier extensions of arbitrary length*, SIAM J. Sci. Comput., 38 (2016), pp. A899–A922.
- [12] R. B. PLATTE, L. N. TREFETHEN, AND A. B. J. KUIJLAARS, *Impossibility of fast stable approximation of analytic functions from equispaced samples*, SIAM Rev., 53 (2011).
- [13] E. A. RAKHMANOV, *Bounds for polynomials with a unit discrete norm*, Ann. of Math. (2), 165 (2007), pp. 55–88.
- [14] L. N. TREFETHEN, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, 2013.
- [15] C. VAN LOAN, *Computational Frameworks for the Fast Fourier Transform*, SIAM, 1992.